

THAT WHICH IS CLAIMED IS:

1. A method of recovering from a failure of a primary routing communication protocol stack which routes
5 communications over connections to a virtual Internet Protocol address (VIPA) and at least one port associated with the VIPA to a plurality of other communication protocol stacks associated by the routing communication protocol stack with the VIPA and the at least one port,
10 the method comprising:

identifying at least one of a plurality of communication protocol stacks other than the primary routing communication protocol stack in a cluster of data processing systems containing the primary routing
15 communication protocol stack as a backup routing communication protocol stack;

storing at the backup routing communication protocol stack an identification of communication protocol stacks associated with the VIPA and the at least one port as
20 candidate target stacks;

receiving messages at the backup routing communication protocol stack identifying ones of the communication protocol stacks having application instances bound to the VIPA and listening on the at least
25 one port as current actual target stacks;

identifying connections to the VIPA and the at least one port for current actual target stacks;

constructing a routing table from the received messages identifying the connections to the VIPA so as to provide routing path information to current actual target
30 stacks associated with the connection; and

routing messages on the connections to the VIPA and the at least one port to the current actual target stacks associated with the connections through the backup

DOVER EDITIONS

routing communication protocol stack utilizing the constructed routing table.

2. A method according to Claim 1, wherein the step
5 of identifying connections to the VIPA comprises the step of receiving messages from the current actual target stacks identifying connections to the VIPA and the at least one port associated with the current actual target stacks.

10

3. A method according to Claim 1, further comprising the step of:

receiving a list of communication protocol stack addresses which are associated with the VIPA and the at
15 least one port; and

wherein the step of storing at the backup routing communication protocol an identification of communication protocol stacks associated with the VIPA and the at least one port as candidate target stacks comprises the step of
20 storing the received list of communication protocol stack addresses.

4. A method according to Claim 3, wherein the list further includes identification of the VIPA and ports
25 associated with the VIPA.

5. A method according to Claim 1, further comprising the step of:

defining, at the backup routing communication
30 protocol stack, an identification of communication protocol stacks which are associated with the VIPA and the at least one port so as to provide a backup list of candidate target stacks which is different from the candidate target stacks of the primary routing
35 communication protocol stacks; and

wherein the step of storing at the backup routing communication protocol an identification of communication protocol stacks associated with the VIPA and the at least one port as candidate target stacks comprises the step of
5 storing the backup list of candidate target stacks.

6. A method according to Claim 5, wherein the communication protocol stacks identified in the backup list of candidate target stacks are selected to reduce
10 disruption of the cluster of data processing systems associated with the communication protocol stacks by failure of the primary routing communication protocol stack.

15 7. A method according to Claim 1, further comprising the step of sending the stored list to the plurality of communication protocol stacks in the cluster of data processing systems.

20 8. A method according to Claim 7, wherein the communication protocol stacks in the cluster of data processing systems carry out the following steps responsive to receiving the list from the backup routing communication protocol stack:

25 transmitting messages to the backup routing communication protocol stack identifying ones of the communication protocol stacks in the cluster of data processing systems which are identified on the list and which have application instances bound to the VIPA and
30 listening on the at least one port;

transmitting messages to the backup routing communication protocol stack identifying connections to the VIPA and the at least one port; and

35 associating the backup routing communication protocol stack with the VIPA and the at least one port so

as to send subsequent messages associated with the VIPA and the at least one port to the backup routing communication protocol stack.

5 9. A method according to Claim 1, further comprising the steps of:

detecting recovery of the primary routing communication protocol stack; and

10 transferring routing of connections to the VIPA and the at least one port back to primary routing communication protocol stack upon detection of recovery of the primary routing communication protocol stack.

15 10. A method according to Claim 9, wherein the steps of detecting recovery and transferring routing comprise the steps of:

20 receiving, at the backup routing communication protocol stack, an identification of communication protocol stacks associated with the VIPA and the at least one port which are being routed by the backup routing communication protocol stack from the primary routing communication protocol stack;

25 associating the primary routing communication protocol stack with the VIPA and the at least one port at communication protocol stacks having connections to the VIPA and the at least one port or identified in the identification from the primary routing communication protocol stack as associated with the VIPA and the at least one port so as to cause subsequent messages associated with the VIPA and the at least one port to be sent to the primary routing communication protocol stack;

30 identifying connections to the VIPA and the at least one port which are routed by the backup communication protocol stack;

constructing a routing table at the primary routing communication protocol stack from the received identification of the connections to the VIPA so as to provide routing path information to current actual target stacks associated with the connection; and

routing messages on the connections to the VIPA and the at least one port to the current actual target stacks associated with the connections through the primary routing communication protocol stack utilizing the constructed routing table.

11. A method of recovering from a failure of a primary routing communication protocol stack which routes messages to a predefined virtual IP address (VIPA) to a predefined set of candidate target communication protocol stacks, the method comprising:

defining a backup routing communication protocol stack which takes ownership of routing functions for the VIPA in the event of failure of the primary routing communication protocol stack;

detecting failure of the primary routing communication protocol stack;

25 determining, at the backup routing communication protocol stack, connections to the VIPA which are associated with the primary routing communication protocol stack;

routing the determined connections to communication protocol stacks associated with the connections through the backup routing communication protocol stack;

30 establishing a second set of candidate target communication protocol stacks which is different from the predefined set of candidate target communication protocol stacks and associated with the backup routing communication protocol stack;

receiving a request for a new connection to the VIPA;

selecting a communication protocol stack from the second set of candidate communication protocol stacks as
5 a destination for the new connection to the VIPA; and

routing messages associated with the new connection to the selected communication protocol stack through the backup routing communication protocol stack.

10 12. A method according to Claim 11, wherein the second set of candidate target communication protocol stacks are selected to reduce disruption of a cluster of data processing systems associated with the communication protocol stacks by failure of the primary routing
15 communication protocol stack.

13. A method according to Claim 11, further comprising the step of sending an identification of the second set of candidate target communication protocol stacks to other communication protocol stacks in a
20 cluster of data processing systems.

14. A method according to Claim 13, wherein the other communication protocol stacks carry out the
25 following steps responsive to receiving the second set of candidate target communication protocol stacks from the backup routing communication protocol stack:

transmitting messages to the backup routing communication protocol stack identifying ones of the
30 other communication protocol stacks which have application instances bound to the VIPA and listening on a port of the VIPA;

transmitting messages to the backup routing communication protocol stack identifying connections to
35 the VIPA; and

associating the backup routing communication protocol stack with the VIPA so as to send subsequent messages associated with the VIPA and the at least one port to the backup routing communication protocol stack.

5

15. A method according to Claim 11, further comprising the steps of:

detecting recovery of the primary routing communication protocol stack; and

10 transferring routing of connections to the VIPA back to primary routing communication protocol stack upon detection of recovery of the primary routing communication protocol stack.

15 16. A method according to Claim 15, wherein the steps of detecting recovery and transferring routing comprise the steps of:

receiving, at the backup routing communication protocol stack, an identification of candidate target communication protocol stacks associated with the VIPA from the primary routing communication protocol stack;

20 associating the primary routing communication protocol stack with the VIPA at communication protocol stacks having connections to the VIPA or identified in the received identification of candidate target communication protocol stacks from the primary routing communication protocol stack so as to cause subsequent messages associated with the VIPA to be sent to the primary routing communication protocol stack;

25 30 identifying connections to the VIPA which are routed by the backup communication protocol stack;

constructing a routing table at the primary routing communication protocol stack from the received identification of the connections to the VIPA so as to provide routing path information for the connections; and

routing messages on the connections to the VIPA to the communication protocol stacks associated with the connections through the primary routing communication protocol stack utilizing the constructed routing table.

5

17. A system for recovering from a failure of a primary routing communication protocol stack which routes communications over connections to a virtual Internet Protocol address (VIPA) and at least one port associated with the VIPA to a plurality of other communication protocol stacks associated by the routing communication protocol stack with the VIPA and the at least one port, comprising:

15 means for identifying at least one of a plurality of communication protocol stacks other than the primary routing communication protocol stack in a cluster of data processing systems containing the primary routing communication protocol stack as a backup routing communication protocol stack;

20 means for storing at the backup routing communication protocol stack an identification of communication protocol stacks associated with the VIPA and the at least one port as candidate target stacks;

25 means for receiving messages at the backup routing communication protocol stack identifying ones of the communication protocol stacks having application instances bound to the VIPA and listening on the at least one port as current actual target stacks;

30 means for identifying connections to the VIPA and the at least one port for current actual target stacks;

means for constructing a routing table from the received messages identifying the connections to the VIPA so as to provide routing path information to current actual target stacks associated with the connection; and

5

means for routing messages on the connections to the VIPA and the at least one port to the current actual target stacks associated with the connections through the backup routing communication protocol stack utilizing the constructed routing table.

10

18. A system according to Claim 17, wherein the means for identifying connections to the VIPA comprises means for receiving messages from the current actual target stacks identifying connections to the VIPA and the at least one port associated with the current actual target stacks.

15

19. A system according to Claim 17, further comprising:

20

means for receiving a list of communication protocol stack addresses which are associated with the VIPA and the at least one port; and

25

wherein the means for storing at the backup routing communication protocol an identification of communication protocol stacks associated with the VIPA and the at least one port as candidate target stacks comprises means for storing the received list of communication protocol stack addresses.

20.

20. A system according to Claim 19, wherein the list further includes identification of the VIPA and ports associated with the VIPA.

30

21. A system according to Claim 17, further comprising:

35

means for defining, at the backup routing communication protocol stack, an identification of communication protocol stacks which are associated with the VIPA and the at least one port so as to provide a

DOVER EDITIONS - 007

backup list of candidate target stacks which is different from the candidate target stacks of the primary routing communication protocol stacks; and

5 wherein the means for storing at the backup routing communication protocol an identification of communication protocol stacks associated with the VIPA and the at least one port as candidate target stacks comprises means for storing the backup list of candidate target stacks.

10 22. A system according to Claim 21, wherein the communication protocol stacks identified in the backup list of candidate target stacks are selected to reduce disruption of the cluster of data processing systems associated with the communication protocol stacks by 15 failure of the primary routing communication protocol stack.

20 23. A system according to Claim 17, further comprising means for sending the stored list to the plurality of communication protocol stacks in the cluster of data processing systems.

25 24. A system according to Claim 23, further comprising:

means for transmitting messages to the backup routing communication protocol stack identifying ones of the communication protocol stacks in the cluster of data processing systems which are identified on the list and which have application instances bound to the VIPA and 30 listening on the at least one port;

means for transmitting messages to the backup routing communication protocol stack identifying connections to the VIPA and the at least one port; and

35 means for associating the backup routing communication protocol stack with the VIPA and the at

● ●

least one port so as to send subsequent messages associated with the VIPA and the at least one port to the backup routing communication protocol stack.

5 25. A system according to Claim 17, further comprising:

means for detecting recovery of the primary routing communication protocol stack; and

10 means for transferring routing of connections to the VIPA and the at least one port back to primary routing communication protocol stack upon detection of recovery of the primary routing communication protocol stack.

15 26. A system according to Claim 25, wherein the means for detecting recovery and transferring routing comprise:

20 means for receiving, at the backup routing communication protocol stack, an identification of communication protocol stacks associated with the VIPA and the at least one port which are being routed by the backup routing communication protocol stack from the primary routing communication protocol stack;

25 means for associating the primary routing communication protocol stack with the VIPA and the at least one port at communication protocol stacks having connections to the VIPA and the at least one port or identified in the identification from the primary routing communication protocol stack as associated with the VIPA and the at least one port so as to cause subsequent messages associated with the VIPA and the at least one port to be sent to the primary routing communication protocol stack;

30 means for identifying connections to the VIPA and the at least one port which are routed by the backup communication protocol stack;

means for constructing a routing table at the primary routing communication protocol stack from the received identification of the connections to the VIPA so as to provide routing path information to current actual target stacks associated with the connection; and

5 means for routing messages on the connections to the VIPA and the at least one port to the current actual target stacks associated with the connections through the primary routing communication protocol stack utilizing
10 the constructed routing table.

27. A system for recovering from a failure of a primary routing communication protocol stack which routes messages to a predefined virtual IP address (VIPA) to a
15 predefined set of candidate target communication protocol stacks, comprising:

means for defining a backup routing communication protocol stack which takes ownership of routing functions for the VIPA in the event of failure of the primary
20 routing communication protocol stack;

means for detecting failure of the primary routing communication protocol stack;

means for determining, at the backup routing communication protocol stack, connections to the VIPA
25 which are associated with the primary routing communication protocol stack;

means for routing the determined connections to communication protocol stacks associated with the connections through the backup routing communication
30 protocol stack;

means for establishing a second set of candidate target communication protocol stacks which is different from the predefined set of candidate target communication protocol stacks and associated with the backup routing
35 communication protocol stack;

means for receiving a request for a new connection to the VIPA;

selecting a communication protocol stack from the second set of candidate communication protocol stacks as
5 a destination for the new connection to the VIPA; and

means for routing messages associated with the new connection to the selected communication protocol stack through the backup routing communication protocol stack.

10 28. A system according to Claim 27, wherein the second set of candidate target communication protocol stacks are selected to reduce disruption of a cluster of data processing systems associated with the communication protocol stacks by failure of the primary routing
15 communication protocol stack.

20 29. A system according to Claim 27, further comprising means for sending an identification of the second set of candidate target communication protocol stacks to other communication protocol stacks in a cluster of data processing systems.

30 30. A system according to Claim 29, further comprising:

25 means for transmitting messages to the backup routing communication protocol stack identifying ones of the other communication protocol stacks which have application instances bound to the VIPA and listening on a port of the VIPA;

30 means for transmitting messages to the backup routing communication protocol stack identifying connections to the VIPA; and

35 means for associating the backup routing communication protocol stack with the VIPA so as to send subsequent messages associated with the VIPA and the at

least one port to the backup routing communication protocol stack.

31. A system according to Claim 27, further
5 comprising:

means for detecting recovery of the primary routing communication protocol stack; and

10 means for transferring routing of connections to the VIPA back to primary routing communication protocol stack upon detection of recovery of the primary routing communication protocol stack.

32. A system according to Claim 31, wherein the
means for detecting recovery and the means for
15 transferring routing comprise:

means for receiving, at the backup routing communication protocol stack, an identification of candidate target communication protocol stacks associated with the VIPA from the primary routing communication protocol stack;

20 means for associating the primary routing communication protocol stack with the VIPA at communication protocol stacks having connections to the VIPA or identified in the received identification of candidate target communication protocol stacks from the primary routing communication protocol stack so as to cause subsequent messages associated with the VIPA to be sent to the primary routing communication protocol stack;

25 means for identifying connections to the VIPA which are routed by the backup communication protocol stack;

30 means for constructing a routing table at the primary routing communication protocol stack from the received identification of the connections to the VIPA so as to provide routing path information for the 35 connections; and

CONFIDENTIAL

means for routing messages on the connections to the VIPA to the communication protocol stacks associated with the connections through the primary routing communication protocol stack utilizing the constructed routing table.

5

33. A computer program product for recovering from a failure of a primary routing communication protocol stack which routes communications over connections to a virtual Internet Protocol address (VIPA) and at least one port 10 associated with the VIPA to a plurality of other communication protocol stacks associated by the routing communication protocol stack with the VIPA and the at least one port, comprising:

a computer readable storage medium having computer readable program code embodied therein, the computer readable program code comprising:

computer readable program code which identifies at least one of a plurality of communication protocol stacks other than the primary routing communication protocol 20 stack in a cluster of data processing systems containing the primary routing communication protocol stack as a backup routing communication protocol stack;

computer readable program code which stores at the backup routing communication protocol stack an identification of communication protocol stacks 25 associated with the VIPA and the at least one port as candidate target stacks;

computer readable program code which receives messages at the backup routing communication protocol 30 stack identifying ones of the communication protocol stacks having application instances bound to the VIPA and listening on the at least one port as current actual target stacks;

00027000000000000000000000000000

computer readable program code which identifies connections to the VIPA and the at least one port for current actual target stacks;

5 computer readable program code which constructs a routing table from the received messages identifying the connections to the VIPA so as to provide routing path information to current actual target stacks associated with the connection; and

10 computer readable program code which routes messages on the connections to the VIPA and the at least one port to the current actual target stacks associated with the connections through the backup routing communication protocol stack utilizing the constructed routing table.

15 34. A computer program product according to Claim 33, wherein the computer readable program code which identifies connections to the VIPA comprises computer readable program code which receives messages from the current actual target stacks identifying connections to the VIPA and the at least one port associated with the current actual target stacks.

20 35. A computer program product according to Claim 33, further comprising:

25 computer readable program code which receives a list of communication protocol stack addresses which are associated with the VIPA and the at least one port; and
wherein the computer readable program code which stores at the backup routing communication protocol an identification of communication protocol stacks
30 associated with the VIPA and the at least one port as candidate target stacks comprises computer readable program code which stores the received list of communication protocol stack addresses.

35

DRAFT - 08/04/00

36. A computer program product according to Claim 35, wherein the list further includes identification of the VIPA and ports associated with the VIPA.

5 37. A computer program product according to Claim 33, further comprising:

10 computer readable program code which defines defining, at the backup routing communication protocol stack, an identification of communication protocol stacks which are associated with the VIPA and the at least one port so as to provide a backup list of candidate target stacks which is different from the candidate target stacks of the primary routing communication protocol stacks; and

15 wherein the computer readable program code which stores at the backup routing communication protocol an identification of communication protocol stacks associated with the VIPA and the at least one port as candidate target stacks comprises computer readable program code which stores the backup list of candidate target stacks.

20 25 38. A computer program product according to Claim 37, wherein the communication protocol stacks identified in the backup list of candidate target stacks are selected to reduce disruption of the cluster of data processing systems associated with the communication protocol stacks by failure of the primary routing communication protocol stack.

30 35 39. A computer program product according to Claim 33, further comprising computer readable program code which sends sending the stored list to the plurality of communication protocol stacks in the cluster of data processing systems.

40. A computer program product according to Claim 39, further comprising:

computer readable program code which transmitts messages to the backup routing communication protocol stack identifying ones of the communication protocol stacks in the cluster of data processing systems which are identified on the list and which have application instances bound to the VIPA and listening on the at least one port;

10 computer readable program code which transmitts messages to the backup routing communication protocol stack identifying connections to the VIPA and the at least one port; and

15 computer readable program code which associates the backup routing communication protocol stack with the VIPA and the at least one port so as to send subsequent messages associated with the VIPA and the at least one port to the backup routing communication protocol stack.

20 41. A computer program product according to Claim 33, further comprising:

computer readable program code which detects recovery of the primary routing communication protocol stack; and

25 computer readable program code which transfers routing of connections to the VIPA and the at least one port back to primary routing communication protocol stack upon detection of recovery of the primary routing communication protocol stack.

30

42. A computer program product according to Claim 41, wherein the computer readable program code which detects recovery and transferring routing comprise:

35 computer readable program code which receives, at the backup routing communication protocol stack, an

identification of communication protocol stacks
associated with the VIPA and the at least one port which
are being routed by the backup routing communication
protocol stack from the primary routing communication
5 protocol stack;

computer readable program code which associates the
primary routing communication protocol stack with the
VIPA and the at least one port at communication protocol
stacks having connections to the VIPA and the at least
10 one port or identified in the identification from the
primary routing communication protocol stack as
associated with the VIPA and the at least one port so as
to cause subsequent messages associated with the VIPA and
the at least one port to be sent to the primary routing
15 communication protocol stack;

computer readable program code which identifies
connections to the VIPA and the at least one port which
are routed by the backup communication protocol stack;

20 computer readable program code which constructs a
routing table at the primary routing communication
protocol stack from the received identification of the
connections to the VIPA so as to provide routing path
information to current actual target stacks associated
with the connection; and

25 computer readable program code which routes messages
on the connections to the VIPA and the at least one port
to the current actual target stacks associated with the
connections through the primary routing communication
protocol stack utilizing the constructed routing table.

30

43. A computer program product for recovering from
a failure of a primary routing communication protocol
stack which routes messages to a predefined virtual IP
address (VIPA) to a predefined set of candidate target
35 communication protocol stacks, comprising:

00000000000000000000000000000000

a computer readable medium having computer readable program code embodied therein, the computer readable program code comprising:

5 computer readable program code which defines a backup routing communication protocol stack which takes ownership of routing functions for the VIPA in the event of failure of the primary routing communication protocol stack;

10 computer readable program code which detects failure of the primary routing communication protocol stack;

15 computer readable program code which determines, at the backup routing communication protocol stack, connections to the VIPA which are associated with the primary routing communication protocol stack;

20 computer readable program code which routes the determined connections to communication protocol stacks associated with the connections through the backup routing communication protocol stack;

25 computer readable program code which establishes a second set of candidate target communication protocol stacks which is different from the predefined set of candidate target communication protocol stacks and associated with the backup routing communication protocol stack;

30 computer readable program code which receives a request for a new connection to the VIPA;

 selecting a communication protocol stack from the second set of candidate communication protocol stacks as a destination for the new connection to the VIPA; and

35 computer readable program code which routes messages associated with the new connection to the selected communication protocol stack through the backup routing communication protocol stack.

DRAFT - DRAFT

44. A computer program product according to Claim
43, wherein the second set of candidate target
communication protocol stacks are selected to reduce
disruption of a cluster of data processing systems
5 associated with the communication protocol stacks by
failure of the primary routing communication protocol
stack.

45. A computer program product according to Claim
10 43, further comprising computer readable program code
which sends an identification of the second set of
candidate target communication protocol stacks to other
communication protocol stacks in a cluster of data
processing systems.
15

46. A computer program product according to Claim
45, further comprising:

computer readable program code which transmitts
messages to the backup routing communication protocol
20 stack identifying ones of the other communication
protocol stacks which have application instances bound to
the VIPA and listening on a port of the VIPA;

computer readable program code which transmitts
messages to the backup routing communication protocol
25 stack identifying connections to the VIPA; and

computer readable program code which associates the
backup routing communication protocol stack with the VIPA
so as to send subsequent messages associated with the
VIPA and the at least one port to the backup routing
30 communication protocol stack.

47. A computer program product according to Claim
43, further comprising:

computer readable program code which detects recovery of the primary routing communication protocol stack; and

5 computer readable program code which transfers routing of connections to the VIPA back to primary routing communication protocol stack upon detection of recovery of the primary routing communication protocol stack.

10 48. A computer program product according to Claim 47, wherein the computer readable program code which detects recovery and the computer readable program code which transfers routing comprise:

15 computer readable program code which receives, at the backup routing communication protocol stack, an identification of candidate target communication protocol stacks associated with the VIPA from the primary routing communication protocol stack;

20 computer readable program code which associates the primary routing communication protocol stack with the VIPA at communication protocol stacks having connections to the VIPA or identified in the received identification of candidate target communication protocol stacks from the primary routing communication protocol stack so as to cause subsequent messages associated with the VIPA to be sent to the primary routing communication protocol stack;

25 computer readable program code which identifies connections to the VIPA which are routed by the backup communication protocol stack;

30 computer readable program code which constructs a routing table at the primary routing communication protocol stack from the received identification of the connections to the VIPA so as to provide routing path information for the connections; and

computer readable program code which routes messages
on the connections to the VIPA to the communication
protocol stacks associated with the connections through
the primary routing communication protocol stack
utilizing the constructed routing table.

5

DOCUMENT # 200-00000000